



US006182136B1

(12) **United States Patent**
Ramanathan et al.

(10) **Patent No.:** **US 6,182,136 B1**
 (45) **Date of Patent:** **Jan. 30, 2001**

(54) **AUTOMATED SERVICE ELEMENTS
 DISCOVERY USING CORE SERVICE
 SPECIFIC DISCOVERY TEMPLATES**

(75) Inventors: **Srinivas Ramanathan**, Sunnyvale;
Deborah L. Caswell, Santa Clara, both
 of CA (US); **Scott S. Neal**; **Marc**
Nijdam, both of Fort Collins, CO (US)
 (73) Assignee: **Hewlett-Packard Company**, Palo Alto,
 CA (US)
 (*) Notice: Under 35 U.S.C. 154(b), the term of this
 patent shall be extended for 0 days.

(21) Appl. No.: **09/149,752**

(22) Filed: **Sep. 8, 1998**

(51) Int. Cl.⁷ **G06F 15/173; G06F 15/16**

(52) U.S. Cl. **709/224; 709/202; 709/220;**
709/223; 709/218

(58) Field of Search **709/200, 201,**
709/202, 217, 221, 224, 226, 229, 227,
249, 223; 370/250, 254

(56) **References Cited**

U.S. PATENT DOCUMENTS

| | | | |
|-----------|-----------|----------------|---------|
| 5,185,860 | 2/1993 | Wu | 709/224 |
| 5,276,789 | 1/1994 | Besaw et al. | 345/449 |
| 5,455,853 | * 10/1995 | Cebulka et al. | 379/207 |
| 5,692,106 | * 11/1997 | Towers et al. | 706/59 |
| 5,778,049 | * 7/1998 | Young et al. | 370/250 |
| 5,796,951 | * 8/1998 | Hammer et al. | 709/223 |


* cited by examiner

Primary Examiner—Zarni Maung
Assistant Examiner—Beatriz Prieto

(57) **ABSTRACT**

A method and system for modeling services available via a network include selecting a core service that is to be modeled, forming a discovery template that is specific to the selected core service, and automatically discovering the elements which cooperate to provide the core service. The discovery template includes instructions for implementing automated techniques for discovering service elements, and preferably services, which are anticipated as being cooperative in executing the core service. The system includes a number of discovery modules for generating outputs indicative of the services and service elements. A discovery engine is responsive to the discovery template to invoke the modules that are identified in the template as being relevant to discovering specified services and service elements. The template also identifies dependencies among the modules, so that the proper sequence of processing can be determined. In one embodiment, the discovery template is organized into sections, with each section (1) being specific to a type of service or service element, (2) specifying at least one discovery routine for identifying the specified type of service or service elements, and (3) specifying dependencies of the identified discovery routine on outputs of other discovery routines. Preferably, each section also includes instructions for configuring the data that is output from the identified discovery routine.

21 Claims, 13 Drawing Sheets






FIG. 1
(PRIOR ART)




FIG. 2




FIG. 3




FIG. 4


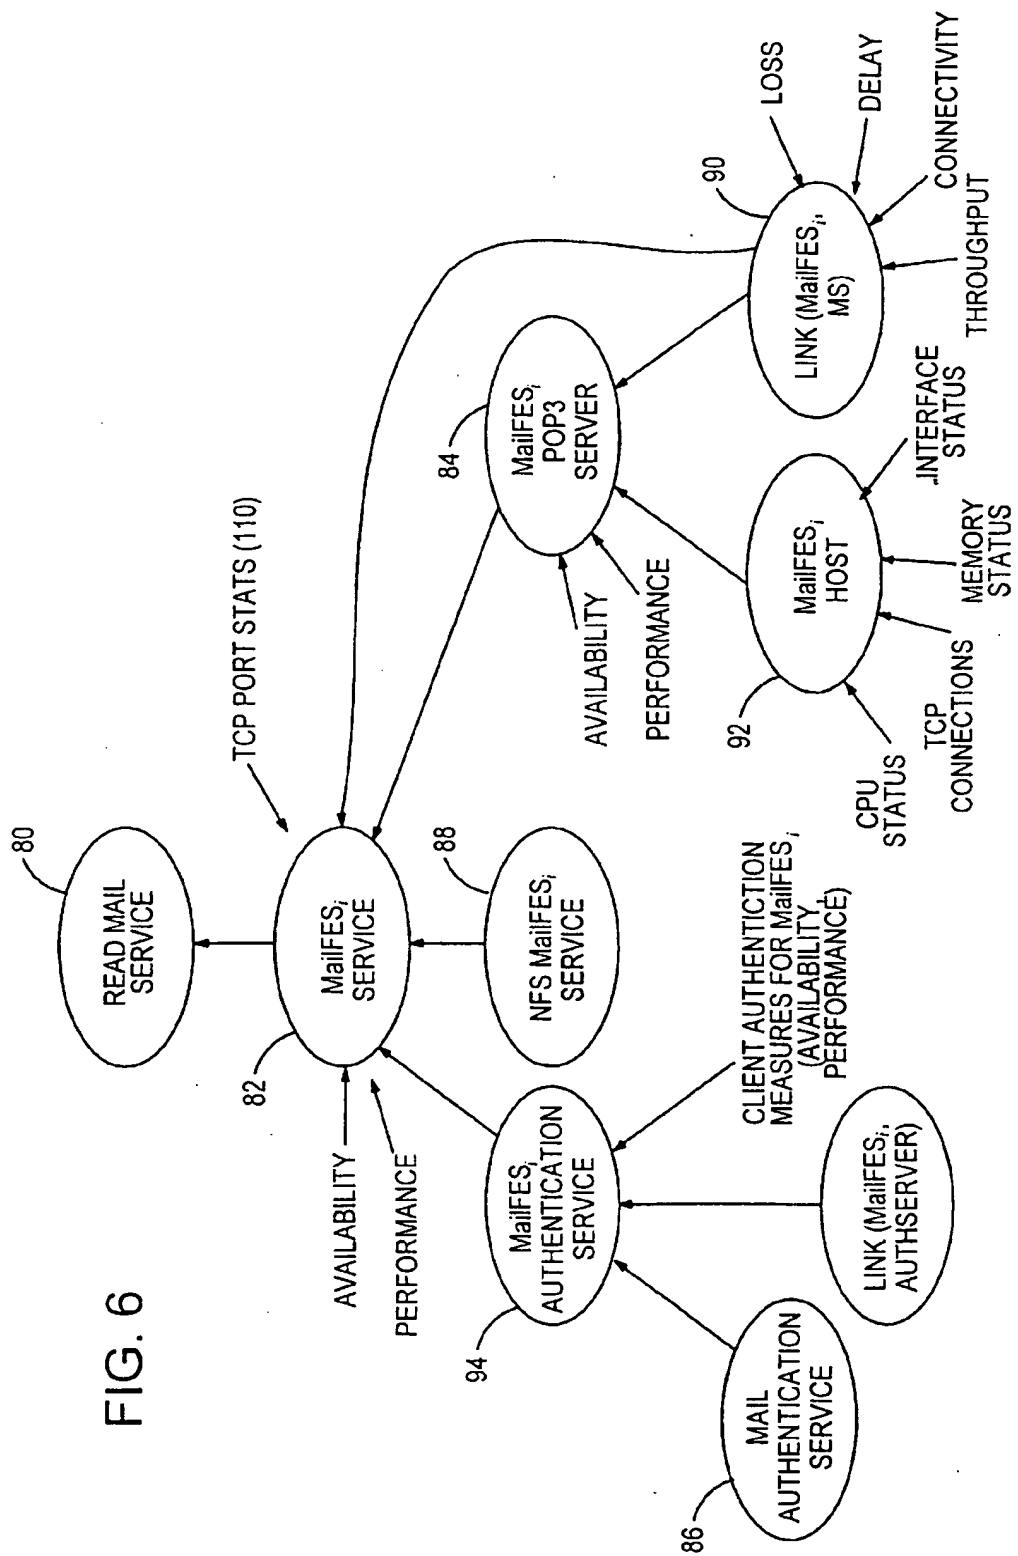


FIG. 5






FIG. 7




FIG. 8




FIG. 9

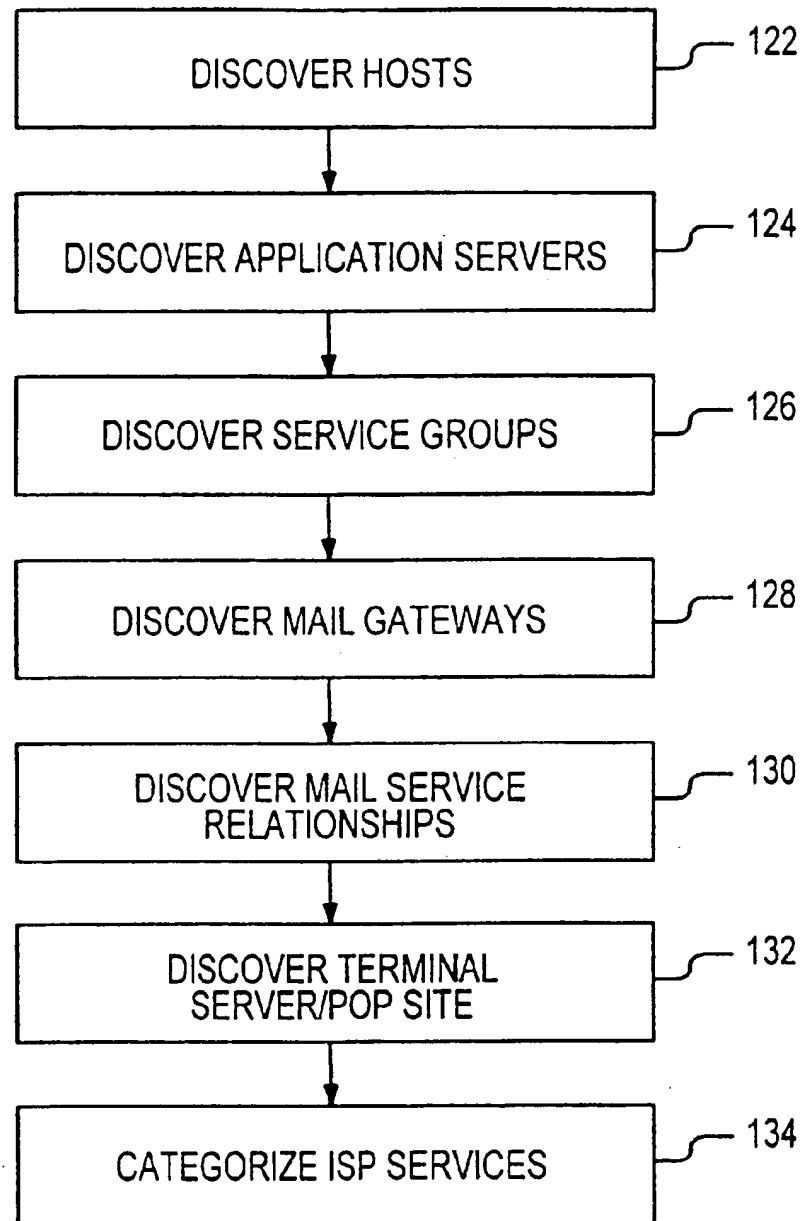


FIG. 10

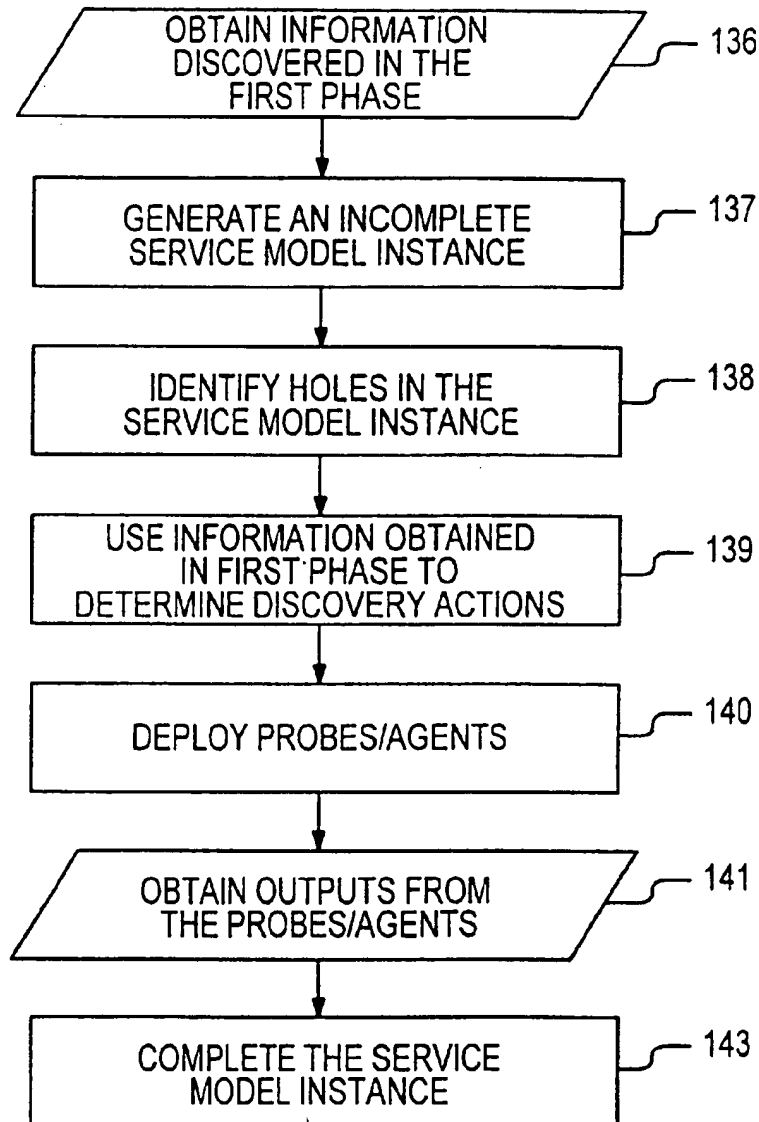


FIG. 11

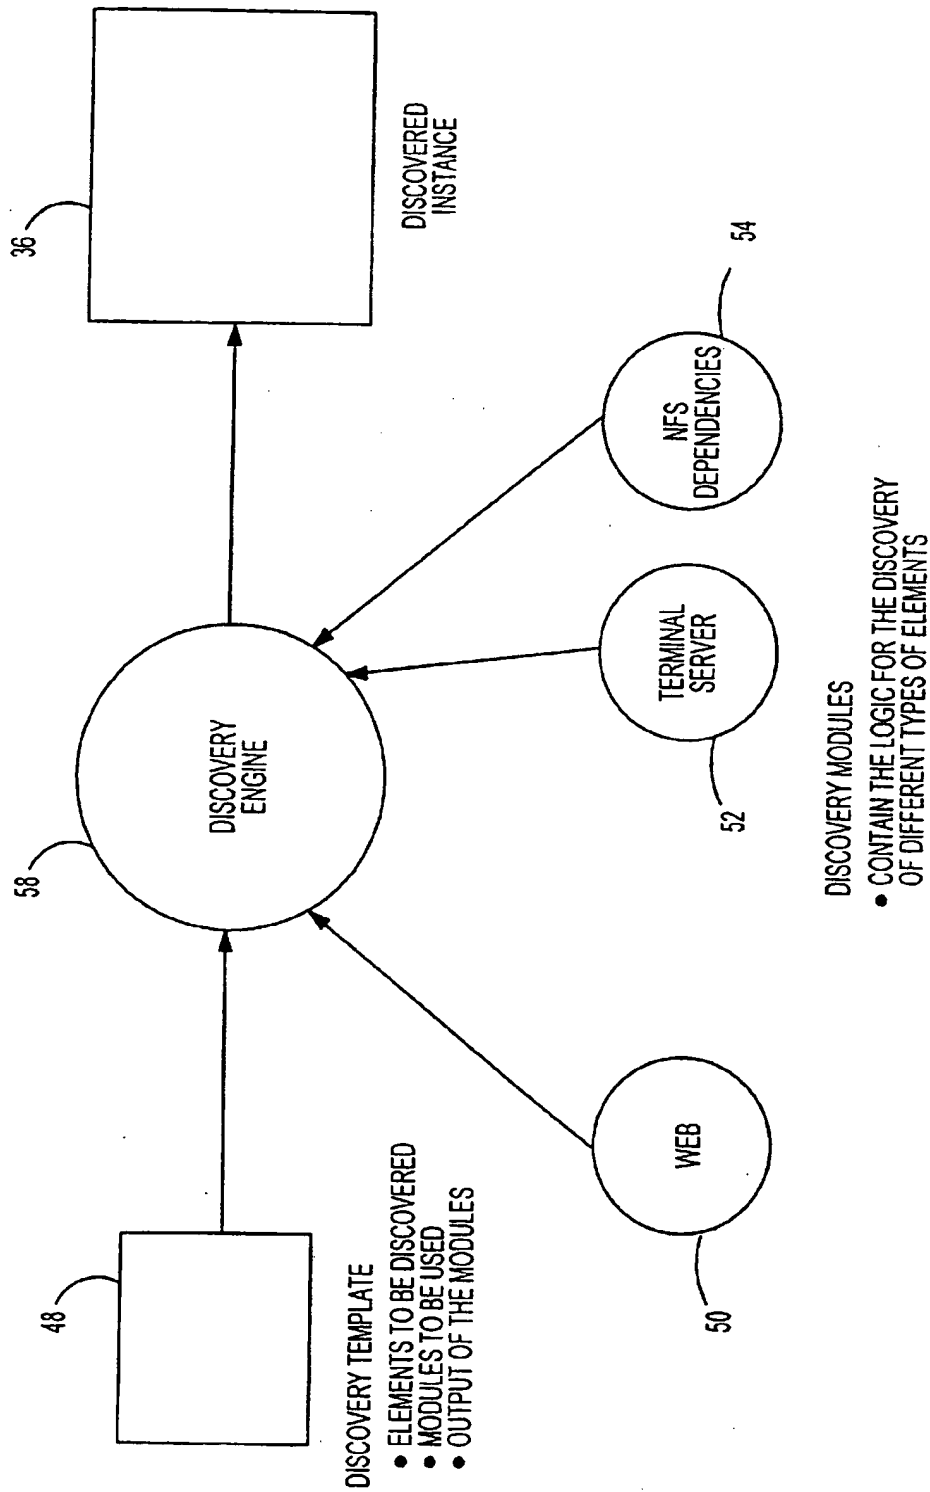


FIG. 12

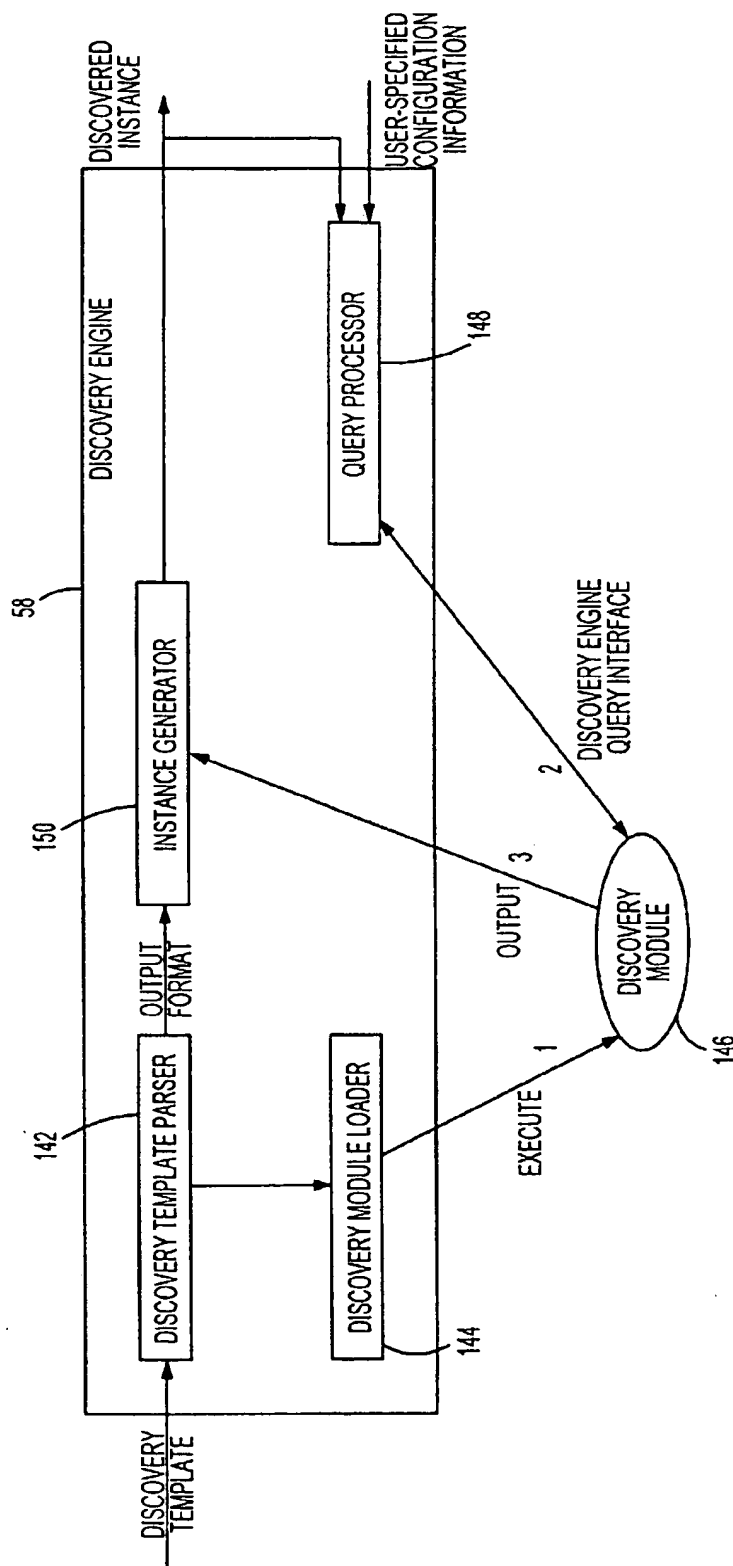


FIG. 13

AUTOMATED SERVICE ELEMENTS DISCOVERY USING CORE SERVICE SPECIFIC DISCOVERY TEMPLATES

TECHNICAL FIELD

The invention relates generally to methods and systems for discovering service elements that enable services via a network, as well as for discovering inter-relationships among the services. Such information may be used for constructing models of services, permitting network personnel to assess the health of the service and to diagnose problems associated with the service.

BACKGROUND ART

Originally, computer networks were designed to have a centralized, network topology. In such a topology, a centralized mainframe computer is accessed by users at computer terminals via network connections. Applications and data are stored at the mainframe computer, but may be accessed by different users. However, a current trend in network design is to provide a topology that enables distributed processing and peer-to-peer communications. Under this topology, network processing power is distributed among a number of network sites that communicate on a peer-to-peer level. Often, there are a number of servers within the network and each server is accessible by a number of clients. Each server may be dedicated to a particular service, but this is not critical. Servers may communicate with one another in providing a service to a client.

Networks vary significantly in scope. A local area network (LAN) is limited to network connectivity among computers that are in close proximity, typically less than one mile. A metropolitan area network (MAN) provides regional connectivity, such as within a major metropolitan area. A wide area network (WAN) links computers located in different geographical areas, such as the computers of a corporation having business campuses in a number of different cities. A global area network (GAN) provides connectivity among computers in various nations. The most popular GAN is the network commonly referred to as the Internet.

The decentralization of computer networks has increased the complexity of tracking network topology. The network components (i.e., "nodes") may be linked in any one of a variety of schemes. The nodes may include servers, hubs, routers, bridges, and the hardware for linking the various components. Systems for determining and graphically displaying the topology of a computer network are known. U.S. Pat. No. 5,276,789 to Besaw et al. and U.S. Pat. No. 5,185,860 to Wu, both of which are assigned to the assignee of the present invention, describe such systems. As described in Besaw et al., the system retrieves a list of nodes and their interconnections from a database which can be manually built by a network administrator or automatically constructed using computer software. The system can be configured to provide any one of three views. An internet view shows nodes and interconnections of different networks. A network view shows the nodes and interconnections of a single network within the internet view. A segment view displays nodes connected within one segment of one of the networks. Selected nodes on the network, called discovery agents, can convey knowledge of the existence of other nodes. The network discovery system queries these discovery agents and obtains the information necessary to form a graphical display of the topology. The discovery agents can be periodically queried to determine if nodes have been added to the network. In a Transmission Controller Protocol/

Internet Protocol (TCP/IP) network, the discovery agents are nodes that respond to queries for an address translation table which translates Internet Protocol (IP) addresses to physical addresses.

The Besaw et al. and Wu systems operate well for graphically displaying hardware components and hardware connections within a network. From this information, a number of conclusions can be drawn regarding the present capabilities and future needs of the network. However, the inter-dependencies of the components in providing a particular service are not apparent from the graphical display that is presented by the system. The complexities of such interdependencies continue to increase in all networks, particularly the Internet. Moreover, these systems are designed in a monolithic manner. This does not allow the management system to be extended to discover and manage new service elements or new services.

Another approach is described by J. L. Hellerstein in an article entitled "A Comparison of Techniques for Diagnosing Performance Problems in Information Systems: Case Study and Analytic Models," *IBM Technical Report*, September, 1994. Hellerstein proposes a measurement navigation graph (MNG) in which network measurements are represented by nodes and the relationships between the measurements are indicated by directed arcs. The relationships among measurements are used to diagnose problems. However, the approach has limitations, since MNGs only represent relationships among measurements. An ISP operator must understand the details of the measurements (when, where, and how each measurement is performed) and their relationships to the different service elements. This understanding is not readily available using the MNG approach. Furthermore, this publication does not explore ways of automatically discovering the relationships among measurements.

The emergence of a variety of new services, such as World Wide Web (WWW) access, electronic commerce, multimedia conferencing, telecommuting, and virtual private network services, has contributed to the growing interest in network-based services. However, the increasing complexity of the services offered by a particular network causes a reduction in the number of experts having the domain knowledge necessary to diagnose and fix problems rapidly. Within the Internet, Internet Service Providers (ISPs) offer their subscribers a number of complex services. An ISP must handle services that involve multiple complex relationships, not only among their service components (e.g., application servers, hosts, and network links), but also within other services. One example is the web service. This service will be described with reference to FIG. 1. Although it may appear to a subscriber of the ISP 10 that the web service is being exclusively provided by a web application server 12, there are various other services and service elements that contribute to the web service. For instance, to access the web server 12, a Domain Name Service (DNS) server 14 is accessed to provide the subscriber with the IP address of the web site. The access route includes one of the Points of Presence (POP) 16, a hub 18, and a router 20. Each POP houses modem banks, telco connections, and terminal servers. A subscriber request is forwarded to and handled by a web server application. The web page or pages being accessed may be stored on a back-end Network File System (NFS) 22, from which it is delivered to the web server on demand. When the subscriber perceives a degradation in the Quality of Service (QoS), the problem may be due to any of the web service components (e.g., the web application server 12, the host machine on which the web application server is

executing, or the network links interconnecting the subscriber to the web server), or may be due to the other infrastructure services on which the web service depends (e.g., DNS or NFS). The ISP system 10 of FIG. 1 is also shown to include an authentication server 24 for performing a subscriber authentication service, a mail server 26 for enabling email service (for login and email access), and front-end and back-end servers 28, 30 and 32 for allowing Usenet access.

Subscribers demand that ISPs offer reliable, predictable services. To meet the expectations of subscribers and to attract new subscribers, ISPs must measure and manage the QoS of their service offerings. This requires a variety of tools that monitor and report on service-level metrics, such as availability and performance of the services, and that provide health reports on the individual service components. Unfortunately, the majority of management systems have not kept pace with the service evolution. Available management systems lack the capability to capture and exploit the inter-relationships that exist among services available in a network environment, such as the Internet. Typically, these management systems discover and manage service elements in isolation. Moreover, these systems are implemented in a monolithic manner and are not easily extensible to discovering and managing new network services and service elements. Adding new discovery and management capabilities to these systems requires extensive redesign and modification of the management system.

Each network is unique in various respects, such as the configuration of servers, the types of application servers, the service offerings, the organizational topology, and the inter-service dependencies. Therefore, in order to accurately understand the operation of the network, specific models must be crafted for the services provided within the network. However, handcrafting models of network services requires an enormous effort on the part of a human expert or group of experts. Furthermore, new Internet services are introduced at a rapid pace. Networking equipment with newer capabilities and faster operation are being introduced rapidly and on an on-going basis. As networks and services evolve, what is needed is a system that is capable of discovering and managing new services and service elements without requiring extensive redesign or modification. Since future services and service elements cannot be predicted with certainty, the system must not be specific to the services and service elements it discovers and manages.

To manage services once they are discovered, what is needed is a method that accommodates the generation of a model of a core service that is available via a network, with the method being responsive to dependencies among services and service elements which are relevant to providing the core service. What is further needed is a system that operates according to the method.

SUMMARY OF THE INVENTION

A method and system for discovering elements which cooperate in the performance of a core service among various services offered by a network include selecting the core service of interest and forming a discovery template that is specific to the selected core service. The discovery template includes data relevant to triggering selected discovery routines for acquiring information that identifies services and service elements that are anticipated as being cooperative in execution of the core service. In the preferred embodiment, the discovery template is specific to the core service, but not specific to any network. When the discovery

template is combined with an engine for driving the discovery routines, information is acquired for mapping dependencies among actual services and actual service elements within the network of interest.

The term "service" as used herein is defined as "a functionality offered by a network to a user to perform a specific set of tasks." Performance of the service involves a number of cooperating elements, such as network routers, servers, applications and the like. Services include application services (such as web and email access) and network services (such as Virtual Private Networks).

The system includes a number of discovery modules for executing the routines designed to detect services and service elements of the network. The system also includes the discovery engine for deploying the modules. The discovery template is accessed by the engine to determine which types of services and service elements are to be discovered and to determine which modules should be deployed. Thus, the discovery template includes identifications of multiple service elements relevant to the core service and includes identifications of discovery modules for detecting the actual services and service elements relevant to the core service. In the preferred embodiment, the discovery template also includes identifications of dependencies among the discovery modules. The discovery template may also include instructions for configuring the instance information that is output from the modules.

The use of the discovery template to provide instance information for modeling the core service preferably occurs without human intervention. Since there are dependencies among the auto-discovery routines executed by the modules, the sequence of executing the routines is typically a concern. In one approach to defining the sequencing of discovery routines, the template is arranged in sections that reflect the dependencies of the discovery modules. That is, a section is processed only after processing of all sections on which it depends. In this approach, the template specifies the order in which discovery is performed. In a second approach, the sequencing is determined by the discovery engine. Again, the discovery template may be divided into sections, but the sequencing of sections is based upon identifications of dependencies by the engine. Sections having no dependencies may be identified and processed first. The discovery engine can then traverse the sections, choosing sections which have not been processed and which have their dependencies determined by earlier processing. This is repeated until all of the sections have been accessed and processed. In a third approach, the discovery modules cooperate to determine the discovery sequencing. For example, a discovery module having a dependency on an output from a second discovery module may register to receive the discovered information from the second module. The dependent module is then notified when the necessary information is available. While this adds complexity to the process, this third approach allows the modules to execute in parallel.

In one embodiment, the network of interest is an Internet Service Provider (ISP) system. The discovery template may be formed to be substantially generic to ISPs, but the instance information received from the discovery modules in response to the discovery template is specific to the ISP of interest. The instance information can then be used to generate a graphical representation of interlinked nodes in which each node represents a service (e.g., mail service and authentication service) or a service element (e.g., host, server and network link) and each interlink indicates an interdependency between two nodes.

An advantage of the invention is that a discovery process occurs automatically in its detection of services and service

elements that cooperate in executing a selected core service. The automated process for discovering services, service elements and dependencies may be executed periodically, so as to discover new services and service elements that may have been introduced into the network environment of interest.

Another advantage is that the system is extensible to discover services or service elements that have been added to a network. When a new service or service element is introduced, a user can add a new discovery module for the service or service element and add the appropriate entries into the discovery template. The discovery engine is then able to discover instances of the new service or service element without any changes or enhancements to the discovery engine. This approach permits third party discovery modules to be easily and quickly introduced, without requiring any changes to the management system.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic view of an exemplary prior art Internet Service Provider (ISP) system.

FIG. 2 is a schematic view of a process for modeling a selected service available via a network.

FIG. 3 is a block diagram of components of a system for modeling a selected service available via a network, such as the ISP of FIG. 1.

FIG. 4 is a block diagram of a management system having components of FIG. 3.

FIG. 5 is an exemplary layout of components for utilizing the Read Mail service of the ISP of FIG. 1.

FIG. 6 is a graph of nodes of a Read Mail service model configured using the system of FIG. 3.

FIG. 7 is an alternative Read Mail service model graph.

FIG. 8 is a schematic view of first phase internal discovery processing using the system of FIG. 3.

FIG. 9 is a schematic view of second phase external discovery processing using the system of FIG. 3.

FIG. 10 is a process flow of steps of the first phase discovery of FIG. 8.

FIG. 11 is a process flow of steps of the second phase discovery of FIG. 9.

FIG. 12 is a block diagram of the operation of the discovery engine of FIG. 4.

FIG. 13 is a block diagram of the discovery engine of FIG. 12.

BEST MODE FOR CARRYING OUT THE INVENTION

The invention relates to the discovery process for enabling automated detection of service elements and/or services that are utilized by a specific network to provide a particular service. One application of the invention is to model the particular service, so that dependencies among services and service elements (e.g., servers, hosts and network links) may be readily determined by network personnel. The discovery process utilizes a service-specific discovery template as a primary means to provide extensibility of an auto-discovery system, since the template drives the performance of the process. Optionally, a service model template may be used in combination with the discovery template to form an instance of the core service offered over the network. However, a service model template is not critical to the practice of the auto-discovery process. It is also noted that while the invention will be described prima-

rily with reference to application to an ISP system, the process has other applications. That is, the process may be used to provide a service model for other providers of services in network environments (e.g., for services within a corporation's network).

FIG. 2 is an overview of the auto-discovery process for generating a service model instance 200 for a particular network 202. In FIG. 2, the rectangles represent data stores and the ellipses represent processing elements. A service-specific discovery template 204 is accessed upon initiation of the auto-discovery process. The discovery template defines the services and service elements that are anticipated as cooperating to provide the core service. Moreover, the discovery template identifies particular discovery tools (i.e., discovery modules and/or agents) that are to be used for each service and service element, as well as any dependencies that a particular discovery tool may have on outputs from other discovery tools. The template preferably includes instructions for configuring the outputs of the discovery tools. A more thorough description of the discovery template will follow, particularly with reference to Table 2, which is an example of a discovery template specification. A more thorough description of the discovery tools 206 will also be provided below, with reference to FIGS. 12 and 13.

The auto-discovery approach is dividable into two phases. In the first phase auto-discovery procedure 208, the information that is available in the domain name system of an ISP is used to discover the existence and the relationships among different services. In order to allow subscribers to access a host using its host name, rather than requiring specification of an IP address that is more difficult to remember, DNS is used to map the host names of the IP addresses for all hosts in the ISP system. Moreover, the exchange of email messages across hosts occurs using the mail exchange records (MX records) maintained by the DNS. In summary, the domain name system stores a wealth of information that is used in the first phase of the auto-discovery process to discover Internet services and relationships among the services. Other mechanisms may be used in this first phase to supplement the information acquired from the domain name system.

Instance information 210 is acquired in the first phase procedure 208. In addition to the identification of services and service elements, the instance information 210 identifies dependencies. There are a number of inter-dependencies of concern to the auto-discovery process. These inter-dependencies include execution dependencies, component dependencies, inter-service dependencies, and organizational dependencies. An execution dependency relates directly to an application server process being executed on a host machine. The types of application servers that are executed on host machines include web, email, news, DNS and NFS. A component dependency occurs in order to ensure scalability and redundancy of a service. For example, a web service may be provided collectively by a number of "front-end servers" (FESs), with round-robin DNS scheduling being used to provide subscribers with a domain name that maps to one of the FESs. ISPs often replicate web, email and news content across a number of servers. The round-robin scheduling balances the load among the servers. The servers are grouped together and assigned a single domain name in the DNS database. When the DNS server receives a request for the domain name, the IP address of one of the servers is acquired in the round-robin scheme.

An inter-service dependency occurs when one service accesses another service for its proper operation. For example, a web service depends on a DNS service to allow

the subscriber to connect the web server host using its IP address, and an NFS service is used to access the web content. As another example, a Read Mail service depends on an authentication service to verify the identity of the subscriber, uses a DNS service to log the subscriber's IP address, and uses an NFS service to access the mail content.

Finally, an organization dependency occurs when there are different ISP operations personnel (e.g., subject matter experts) who are responsible for different services and service elements. For example, an ISP may have a first supervisor managing the web service, a second supervisor managing DNS, and a third supervisor managing NFS. Operational responsibilities may be delegated based upon the geographical location of the service elements. Since the precise organization structure may vary from one ISP to another, the auto-discovery mechanism provides a means by which it can be quickly customized for a particular ISP system.

The first phase auto-discovery procedure 208 provides discovered instance information 210 that identifies most of the execution, component and organization dependencies, as well as some of the inter-service dependencies. A partial service model instance generator 212 is then used to provide a partial service model instance 214. Optionally, configuration data 216 may be used to allow an operator to customize a service model instance 200. Using a configuration interface which will be described with reference to FIG. 4, the operator can specify categorization criteria for services and service elements. Thus, the service model instance 200 can be configured to meet the specific needs of the operator.

In a second phase auto-discovery procedure 218, an internal view of the network 202 is acquired. The internal discovery of the second phase is intended to fill any "holes" in the partial service model 214, and particularly focuses on identifying inter-service dependencies. There are two basic approaches to the internal discovery of the second phase. One approach is the use of network probes, which are implemented in software and are installed at strategic locations on the network to acquire information from headers of packet transmissions. By processing the headers of packets, a software probe can deduce many of the relationships that exist among servers. The second basic approach is to use special-purpose discovery agents that are installed on the ISP hosts to discover relationships among services. Rather than examining headers of packet transmissions, the special-purpose agents use a number of operating systems and application-specific mechanisms (such as processing service configuration information and application-dependent monitoring tables) to discover inter-service dependencies.

The inter-service dependency information 220 from the second phase auto-discovery procedure 218 is combined with the partial service model instance 214 using a second service model instance generator 222. The output of the second service model instance generator is the completed service model instance 200.

The invention will be described primarily in its preferred embodiment of using the discovery template to detect both services and service elements to form a service model. However, a discovery template in accordance with the invention may be restricted to detecting service elements and the discovered instance information may be utilized for purposes other than forming a model of the type to be described below.

ONE EMBODIMENT FOR USING THE INSTANCE INFORMATION

As previously noted, the use of the discovery template 204 is optionally combined with a service model template in

the process of generating the service model instance 200. An overview of the template-driven procedure is illustrated in FIG. 3. The service model template is a generic specification of the service topology and measurement topology for the service of interest (e.g., Read Mail service). Depending on the service being modeled and the service elements that are likely to be involved, the template defines nodes of various types (e.g., hosts, servers, network links, and services) and their associated measurements. Moreover, the template indicates the dependencies among the nodes, such as the dependency of the service on other services (e.g., the Read Mail service which refers to a subscriber accessing his/her mailbox depends on the authentication and NFS services). In the preferred embodiment, the template also includes default state computation rules for specified nodes, so that the state (i.e., "health") of a node can be computed based upon measurements associated with the node and upon states of dependencies of the node.

The service model template 34 is not specific to any ISP system in the sense that it does not specifically reference any hosts, servers, network links, service groups, or other services or service elements in the ISP system. The template may be considered as a "lifeless tree." There is no association between nodes in the service model template and running elements, such as hosts and servers. However, the information contained in the service template for a node may include the element type, the element dependencies, the measurement definitions (e.g., agent to run, the format of the run string, the number and type of parameters, and the format of the output), default state computation rules, default thresholds, default baselines, and default alarm generation and correlation rules.

In database terminology, the service model template 34 is the schema. On the other hand, an instance defines the records in the database and the values of the static information. In FIG. 3, the discovered instance 36 is determined using auto-discovery, as will be explained fully below. Information regarding the services and service elements (e.g., servers, hosts and links) that exist in the ISP system or other service provider systems may be auto-discovered. The store representing auto-discovered information shall be referred to as the auto-discovered instance 36.

The service model creation engine 38 of FIG. 3 is used to generate a service model instance 40 based on the service model template 34 and auto-discovered instance 36. Ideally, all of the discovered information is available prior to instantiation. However, in many cases, discovery has to be performed in multiple phases, such as the two outlined above. In such cases, instantiation may occur partially after each phase of discovery. The main advantage of providing a partial service model instance is that the partially completed service model can provide a guide to identify the additional information needed in subsequent phases of discovery. The service model creation engine 38 encompasses the functions of the partial service model instance generator 212 and the second service model instance generator 222 of FIG. 2. Since new elements and services may be added to the ISP system over time, the service model instantiation process has to be repeated periodically to augment the initially created service model instance.

Unlike the service model template 34, the service model instance 40 is specific to an ISP system. The process of constructing the service model instance using the service model template and the auto-discovered instance 36 is referred to as the "instantiation" of the service model. As previously noted, the relationship between the service model template and the auto-discovered instance is analogous to the relationship between the schema and records in a database.

The service model instance 40 maps services and service elements that exist in a particular ISP system with nodes in the service model template. In doing so, the service model instance also specifies a set of measurements that must be made against each of the services and service elements in the specific ISP system.

The service model instance 40 may be used by a view generator 42. In the preferred embodiment, the service model instance is represented as a graph of the nodes and edges that identify dependencies of the nodes. Different management functions of an ISP will benefit from viewing different subsets of nodes and edges of the service model instance 40. Thus, the view generator 42 may be used by operations personnel to provide an "operations view" of only the services and service elements that fall in a particular domain of control of the personnel. On the other hand, a "customer support view" may provide an end-to-end view for the customer support personnel of the ISP. A "planning view" may be used to graphically present information about usage and performance trends, so that future requirements can be ascertained.

Even when different management functions are interested in the same subset of nodes and edges of the service model instance 40, there may be different interests with regard to rules to be used for state computations. For instances, a management application that visually depicts a color-coded hierarchical representation of the service model instance to operations personnel may require that hierarchical state propagation rules be employed in the service model instance. In this manner, viewing the state of the top-level nodes of the service model, an ISP operator can determine whether any problem has been detected in the ISP system. In contrast, an application that is responsible for automatically detecting, diagnosing and reporting the root-causes of problems in the ISP system may prefer not to use a hierarchical state propagation rule, since it is interested in the state of each node of the service model independently of the state of other nodes.

The view generator 42 may be used to define the subset of nodes of a service model instance that are of interest to a management application, as well as the method for evaluating the states of the nodes of interest. The measurements that are associated with the nodes in the service model instance are common across the different views. However, the rules that apply to computing the state of a particular node based upon relevant measurements and the states of dependent nodes may be different for the different views.

A measurement agent configurator 44 may be used to extract information from the service model instance 40, when the information is relevant to scheduling tests and retrieving test results. At the completion of the service model instance 40, static configuration information as to how the tests of various elements are to be run may be merged with default algorithm descriptions into a service model instance file. State computation rules, thresholds and alarm correlation state defaults may be overridden using the measurement agent configurator. Specific measurements may be enabled or disabled. A final measurement agent specification 46 is generated by the configurator for the specific ISP. The measurements that are identified in the specification 46 are executed using software agents 45. The results of the measurements are analyzed and used by a service model manager 47 to compute the states of different nodes in the service model.

OVERVIEW OF INSTANTIATION

FIG. 4 represents an overview of the process for generating the discovered instance that is used to form the service

model instance 40 of FIG. 3. As previously noted, a requirement of the instantiation of a service model is the generation of the service model template for the particular service or services of interest. The template may be handcrafted, preferably by a domain expert of the service. The template includes a specification of the elements involved in providing the service, such as servers and network links, and the dependencies of the service on other services, such as the dependency of Read Mail service on the authentication and NFS services. As will be explained more fully below, the discovery process includes designation of a discovery template 48. The discovery template specifies the types of services and the service elements to be discovered. The discovery template also includes specifications of discovery modules 50, 52 and 54 to be invoked in the discovery of the specified services and service elements.

The service model template 34, the discovery template 48, and the discovery modules 50, 52 and 54 are utilized within a management system 56 for forming the service model instance 40. Another key component of the management system is a discovery engine 58 that processes information contained within the discovery template 48. The discovery engine invokes the appropriate discovery modules 50-54, interprets the outputs of the invoked modules, and stores the outputs in memory (or in a database) as the discovered instance 36, which is made accessible to the service model creation engine 38. The discovery engine 58 supports a configuration interface 60 that allows ISP operations personnel to control and customize the discovery process. Through the configuration interface 60, an ISP operator can restrict the discovery to specified IP address ranges or host name patterns. Furthermore, the operator can specify naming conventions used by the ISP (e.g., for naming terminal servers and POP sites), which are used by some of the discovery modules 50-54.

The configuration interface 60 serves as a way for an ISP operator to quickly customize the service model instance 40 that is generated by the process. Using the configuration interface, the operator can also specify categorization criteria for services and service elements. For instance, all the mail services could fall in one category, while the DNS servers could fall in another. The categories assigned to services and service elements can represent the organizational structure of the ISP, the geographical location of the servers offering the services (e.g., servers in San Francisco fall in one category, while servers in New York fall in another), or differences in business functions of the service (e.g., web servers that an ISP is hosting on behalf of business customers, as opposed to local web servers that the ISP offers for providing access to the dial-up subscribers).

In the preferred embodiment, the configuration interface 60 is implemented using a graphical user interface (GUI) at an operator computing station. An example of a configuration specification that is entered using the interface 60 is shown in Table 1.

TABLE 1

| | | |
|---|---------------|---------------------------|
| Hosts | exclude ip | 10.1.204.1-10.1.205.1 |
| #exclude all hosts with IP addresses in this range. These hosts represent subscriber home PCs | | |
| WebServers | category name | *.com.fakeisp.net WebHost |
| #servers with names of the form *.com.fakeisp.net are Web hosting servers | | |

TABLE 1-continued

TerminalServers extract name max{0-9}{3}<PopSite>{0-9}.
fakeisp.net
#Terminal servers must match the naming pattern above. Extract the POP
site name from the terminal server's name

The discovery modules 50-54 obtain the configuration criteria from the discovery engine 58. The modules execute the appropriate discovery techniques, and as part of their outputs, record the categories to which the different services and service elements belong. The category information is stored at the discovery instance 36 and is interpreted by the service model creation engine 38 in a manner that is reflective of the service model template 34.

The following sections provide details regarding implementation of the process described with reference to FIGS. 2 and 3. Since the service models template 34 depends on the syntax specified in the discovery template 48, the discovery template specification is considered first. All the templates and instances presented as examples use the INI file format that is commonly used to represent the content of system files in Microsoft Windows-based personal computer systems. However, the process may be implemented using other specification and schema definition technologies (e.g., the Common Information Model specified by the Desktop Management Task Force). Per the INI format, a template or an instance is organized as different sections, each of which is specified by a unique stream enclosed within a pair of square brackets ("[<section name>]").

DISCOVERY TEMPLATE SPECIFICATION

An example discovery template specification is shown in Table 2. Each section of the discovery template defines a specific service or service element. The first four sections represent templates for discovery of external name servers, hosts, Mail SMTP servers, and Mail POP3 servers. The "module" variable specification in each section identifies the discovery module 50-54 of FIG. 4 that is to be used for the particular service or service element. The "arguments" variable represents arguments that are to be passed by the discovery engine 58 to the discovery module during deployment. The "outputs" variable defines the number and names of the columns in the output of the discovery module. The "instanceKey" variable denotes the index that is to be used to access the discovery instance 36 corresponding to each row of output generated by the discovery module. The name or names specified within angled brackets (<>) on the right side of the instanceKey assignment must correspond to one of the column names specified in the output assignment. Finally, the "dependencies" variable indicates the dependencies that a discovery module has on other modules. The discovery engine may use this information to select a sequence in which it processes the discovery templates.

TABLE 2

[ExternalNameServers]
discovery-module=DiscoverExternalNameServers.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=ipAddress, hostName, domainName, category
discovery-instanceKey=<ipAddress>:DNS
discovery-dependencies=
[Hosts]
discovery-module=DiscoverHosts.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl

TABLE 2-continued

discovery-outputs=ipAddress, hostName, state, category
discovery-instanceKey=<ipAddress>:Host
discovery-dependencies=ExternalNameServers
[SMTPServers]
discovery-module=DiscoverSmtpServers.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=ipAddress, hostName, category
discovery-instanceKey=<ipAddress>:Smtp_Mail
discovery-dependencies=Hosts,ExternalNameServers
[POP3Servers]
discovery-module=DiscoverPop3Servers.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=ipAddress, hostName, relatedSmtpServer, category
discovery-instanceKey=<ipAddress>:Pop3_Mail
discovery-dependencies=Hosts,ExternalNameServers
[HTTPServers]
discovery-module=DiscoverHttpServers.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=ipAddress, hostName, serverType, category
discovery-instanceKey=<ipAddress>:Web
discovery-dependencies=Hosts
[NFSServers]
discovery-module=DiscoverNFSServers.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=ipAddress, hostName, category
discovery-instanceKey=<ipAddress>:NFS
discovery-dependencies=Hosts
[Host-Groups]
discovery-module=DiscoverHostGroups.class
discovery-arguments=url http://ism-pc2/scripts/discEngineAPI.pl
discovery-outputs=groupName, groupComponentsList, category
discovery-instanceKey=<groupName>:HostGroup
discovery-dependencies=ExternalNameServers
[WebServiceGroups]
hostName=mailfes21.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
[10.137.196.54:Pop3_Mail]
ipAddress=10.137.196.54
hostName=mailfes22.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
[10.137.196.56:Pop3_Mail]
ipAddress=10.137.196.56
hostName=mailfes23.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
[10.137.196.58:Web]
ipAddress=10.137.196.58
hostName=www.fakeisp.net
serverType=NCSA/1.5
category=InternalWeb
[10.137.196.69:Web]
ipAddress=10.137.196.69
hostName=www.xyz.com.fakeisp.net
serverType=Apache/1.2
category=WebHost
[10.137.196.70:Web]
ipAddress=10.137.196.70
hostName=www.abc.com.fakeisp.net
serverType=Apache/1.2
category=WebHost
[10.174.173.23:NFS]
ipAddress=10.174.173.23
hostName=dns1.fakeisp.net
[pop3.fakeisp.net:HostGroup]
groupName=pop3.fakeisp.net
groupComponentsList=10.137.196.52:10.137.196.54:10.137.196.56
[pop3.fakeisp.net:Pop3_MailServiceGroup]
serviceGrpName=pop3.fakeisp.net
serviceGrpComponentsList=10.137.196.52:10.137.196.54:10.137.196.56
[ExternalDnsServer:Category]
categoryName=ExternalDnsServer
[InternalWebServer:Category]
categoryName=InternalWeb
[WebHostedServer:Category]
categoryName=WebHost

Because the ISP environment has a heterogeneous set of elements (e.g., host nodes, routers, application servers, services and inter-service dependencies), sections of the discovery template must be processed in an order in which elements having no dependencies are considered first. Sections having dependencies can be processed after the dependencies have been completed. There are at least three approaches to ordering the processing of sections. One approach is to order the sections in the template to reflect the dependencies among sections. Thus, a section appears in the template only after the appearance of all sections on which it depends. The discovery engine 58 can then process the sections in order. Another approach is to allow the discovery engine to dictate the sequence of discovery. By considering the values of the dependencies variable within the sections of the discovery template, the discovery engine can determine the order in which the sections must be processed. Sections of the template having no dependencies are processed first. After all such sections are processed, the discovery engine iterates through the list of template sections, choosing sections which have not been processed and which have their dependencies determined by earlier processing. This procedure is repeated until all of the sections have been processed. In the third approach, the sequencing is driven by the discovery modules 50-54 themselves. In this embodiment, the discovery engine processes the template once, invoking the discovery modules simultaneously. The discovery modules determine when the different elements of the ISP system are discovered. When a new instance is detected by a discovery module, it forwards the results to the discovery engine. Based on the dependencies on the discovery module, as specified in the discovery template, the engine forwards the results to other discovery modules for which the results are relevant. The availability of the new results may trigger discovery by other modules. This procedure is repeated until all of the sections have been fully processed. In an alternative implementation, the discovery modules can communicate with one another without involving the discovery engine.

DISCOVERED INSTANCE SPECIFICATION

Table 3 is a portion of an example of the discovered instance 36 of FIG. 4. Section names in the discovered instance correspond to the instanceKey variable specifications and the discovery template of Table 2. For each of the output variables in the discovery template, there is an assignment in the discovered instance of Table 3. The ISP system being discovered in this example is assumed to have the domain name fakeisp.net.

TABLE 3

```
[10.174.173.23:DNS]
ipAddress=10.174.173.23
hostName=dns1.fakeisp.net
domainName=fakeisp.net
category=ExternalDnsServer
[10.174.173.23:Host]
ipAddress=10.174.173.23
hostName=dns1.fakeisp.net
state=Alive
category=
[10.137.196.52:Host]
ipAddress=10.137.196.52
hostName=mailfes21.fakeisp.net:smtp.fakeisp.net
state=Alive
category=
[10.137.196.54:Host]
ipAddress=10.137.196.54
```

TABLE 3-continued

```
hostName=mailfes22.fakeisp.net
state=Alive
category=
[10.137.196.56:Host]
ipAddress=10.137.196.56
hostName=mailfes23.fakeisp.net
state=Alive
category=
[10.137.196.58:Host]
ipAddress=10.137.196.58
hostName=www.fakeisp.net
state=Alive
category=
[10.137.196.69:Host]
ipAddress=10.137.196.69
hostName=www.xyz.com.fakeisp.net
state=Alive
category=
[10.137.196.70:Host]
ipAddress=10.137.196.70
hostName=www.abc.com.fakeisp.net
state=Alive
category=
[10.137.196.52:Smtp_Mail]
ipAddress=10.137.196.52
hostName=mailfes21.fakeisp.net:smtp.fakeisp.net
category=ExternalSmtpServer
[10.137.196.52:Pop3_Mail]
ipAddress=10.137.196.52
hostName=mailfes21.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
[10.137.196.54:Pop3_Mail]
ipAddress=10.137.196.54
hostName=mailfes22.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
[10.137.196.56:Pop3_Mail]
ipAddress=10.137.196.56
hostName=mailfes23.fakeisp.net
relatedSmtpServer=smtp.fakeisp.net
category=
```

Next, a portion of a service model template specification is presented in Table 4 as an example service model template 34. The service model template contains the intelligence to map the discovered instance into the service model nodes. The discovery modules (and hence the discovered instance) are designed independently of the service model instance being generated. This enables the same discovered instance to be used in the generation of different service models (e.g., for different services). Each section in the service model template represents a type of node in the service model instance 40 and contains a series of instructions for creating a node in the service model instance. The service model creation engine 38 processes sections of the service model template, one at a time, attempting to match the template with elements in the discovered instance 36. Each element in the discovered instance corresponds to a section of the discovered instance specification. Lines beginning with a ";" represent comments. These lines are ignored by the service model creation engine 38 when it processes the service model template.

TABLE 4

```
[SM-Host]
;Host Node in the service model
match={"*.Host"}
;for each discovered Host
instancekey=<ipAddress>-SM-Host
measurements=TCP-Connection Rate(<ipAddress>),
```

TABLE 4-continued

```

VMstat(<ipAddress>).IFstat(<ipAddress>)
;these are measurements of the host
;next copy its attributes to the SM node
hostname=<hostName>
ipAddress=<ipAddress>
state=<state>
category=<category>
[SM-Web]
;SM node for a web server
match=[*:Web]
;match all discovered Web server instances
instanceKey=<ipAddress>:SM-Web
components=<ipAddress>:SM-Host,<ipAddress>:mslP:SM-Link
measurements=HTTP-TCPConnectionTime(<ipAddress>)
;next copy all the attributes from the server discovered instance
stateComputationRule=measurementsOnly
serverType=<serverType>
hostName=<hostName>
ipAddress=<ipAddress>
category=append(<category>,<ipAddress>:SM-Host?<category>)
[SM-WebService]
;a web service node
match=[*:SM-Web]
;there is a web service node corresponding to each web server node
instanceKey=<ipAddress>:SM-WebService
measurements=HTTP-Availability(<ipAddress>), HTTP-Total
ResponseTime(<ipAddress>), HTTP-DnsTime(<ipAddress>)
components=<ipAddress>:SM-Web,<ipAddress>:Web:SM-NFS,
<ipAddress>:Web:SM-DNS
;NFS and DNS service dependencies will be determined by phase 2
discovery category=<category>
[SM-WebServiceGroup]
match=[*:WebServiceGroup]
instanceKey=<serviceGrpName>:SM-WebServiceGroup
components=list(<serviceGrpComponentsList>):SM-WebService
category=<category>
category=append(list(<serviceGrpComponentsList>):SM-WebService?
<category>)
[SM-TopLevel-Web]
instanceKey=Web:SM-TopLevelWeb
components=*:SM-WebService, *:SM-WebServiceGroup
;components are all web services and all web service groups
[SM-DNS]
match=[*:DNS]
instanceKey=<ipAddress>:SM-DNS
components=<ipAddress>:mslP:SM-Link
measurements=DNS-Availability(<ipAddress>),
DNS-CacheHitResponseTime(<ipAddress>)
stateComputationRule=measurementsOnly
ipAddress=<ipAddress>
hostName=<hostName>
domainName=<domainName>
category=<category>
[SM-NFS]
match=[*:NFS]
instanceKey=<ipAddress>:SM-NFS
components=<ipAddress>:SM-Host
measurements=NFS-TotalCalls(<ipAddress>), NFS-DupReqs
(<ipAddress>), NFS-TimeOutPercent(<ipAddress>)
stateComputationRule=measurementsOnly
ipAddress=<ipAddress>
hostName=<hostName>
[SM-Pop3_Mail]
match=[*:Pop3_Mail]
instanceKey=<ipAddress>:SM-POP3_Mail
components=<ipAddress>:SM-Host
measurements=POP3-TCPConnectionTime(<ipAddress>)
hostName=<hostName>
category=<category>
relatedSmtpServer=<relatedSmtpServer>
ipAddress=<ipAddress>
[SM-ReadMailService]
match=[*:SM-Pop3_Mail]
instanceKey=<ipAddress>:SM-ReadMailService
measurements=POP3-Availability(<ipAddress>),
POP3-TotalResponseTime(<ipAddress>),
POP3-AuthenticationTime(<ipAddress>)
stateComputationRule=default
components=<ipAddress>:SM-Pop3_Mail,<ipAddress>:Pop3_Mail:

```

TABLE 4-continued

```

SM-NFS>,<ipAddress>:Pop3_Mail:SM-Auth>
category=<category>
5 [SM-ReadMailServiceGroup]
match=[*:Pop3MailServiceGroup]
instanceKey=<serviceGrpName>:SM-ReadMailServiceGroup
components=list(<serviceGrpComponentsList>):SM-ReadMailService
category=<category>
[SM-TopLevel-ReadMail]
10 instanceKey=ReadMail:SM-TopLevel-ReadMail
components=*:SM-ReadMailService, *:SM-ReadMailServiceGroup
[SM-Category]
match=[*:Category]
instanceKey=<categoryName>:SM-Category
components=*:SM-*?category=<categoryName>
15

```

Most sections of the service model template 34 begin with a "match" criteria. The match criteria for a section of the service model template specifies the discovery instances that are relevant to the section under consideration. For instance, the match criteria corresponding to the host node's specification (SM-Host) in the discovery template indicates that the corresponding discovered instances are those of type Host. The match criteria is specified as a regular expression that is matched against section names (instance keys) of the discovered instance 36. For each object (section) in the discovered instance that matches the regular expression specified in the match criteria, a corresponding node is instantiated in the service model instance 40.

Each section of the service model template 34 can match discovered instances of at least one type. When a discovered instance satisfies the match criteria specified in a section of the service model template 34, any of the attributes of the discovered instance can be referred to in the subsequent instructions of the service model template's section. The absence of a match criteria in the specification of a section of the service model template indicates that there is only one instance of that type for the particular ISP.

The instanceKey variable in Table 4 denotes the key that is to be used to refer to a service model node that is instantiated by the section of the template under consideration. The attributes enclosed within the angled brackets ("<>") must be one of the attributes of the elements of the discovered instance for which there is a reference by the match criteria.

The "components" instruction specifies the parent-child relationship in a service model instance. Various types of dependencies (i.e., execution dependencies, component dependencies, inter-service dependencies and organizational dependencies) are captured by this specification. The components list specified must make reference to the node in the service model instance 40 that is to be generated from the service model template 34. The components list refers to specific nodes, all nodes of a specific type, and all nodes of a different specific type that have a specific attribute value. Sections of the service model template that refer to leaf nodes of the service model instance do not have component specifications.

The "measurements" instruction specifies a list of measurements that must be targeted at the corresponding node in the service model instance 40. By processing the measurement specifications of nodes in the service model instance, the measurement agent configurator 44 of FIG. 3 can determine the agents that must be scheduled for execution against each element of the discovered instance 36. It should be noted that not all nodes in the service model instance have measurement specifications.

The "StateComputationRule" instruction covers how the states of the corresponding nodes in the service model

17

instance 40 are computed. By default, the state of a node in the service model instance is determined based on the states of all of the measurements associated with a node and the states of all of its dependencies (children nodes) in the service model instance. The service model creation engine 38 may support additional state computation policies. For example, a "measurementsOnly" policy indicates that only the states of the measurements associated with a node must be taken into account in determining the state of that node.

Regarding attribute value settings, each service model node may derive attributes from the discovered instance 36 to which it refers. The service model template syntax also allows for hierarchical aggregation of attributes. This is demonstrated in the append construct used for defining category attributes for the service model nodes.

SERVICE MODEL CREATION ENGINE

The service model creation engine 38 incorporates the logic for processing a service model template 34 with the discovered instance 38 to generate the service model instance 40. There are alternatives with regard to the order in which the engine 38 processes the service model template. In a sequential processing approach, it is assumed that the service model template was constructed such that the sections of the service model template are in an order in which they need to be processed. The engine can then process the sections sequentially. The sequential processing enables simplification of the service model creation engine. However, this approach burdens the template developer with a requirement of manually determining the placement of each section in the template based upon the order of processing. Moreover, since processing typically starts from the host nodes, this approach may result in a number of service model host nodes that do not have additional nodes above them when a service model instance graph is generated in a manner to be described below. To avoid such "orphaned" nodes, the created service model instance must be further processed.

In the alternative hierarchical processing, a service model creation engine 38 can use the "components" specifications in the different sections of the service model template 34 to determine the order for processing the sections of the template. Since the components list specifies the dependencies of a node, before processing a section, all of the sections corresponding to each of the nodes types in the components list must be processed. Based on this rule, sections of the templates which have no specified components are processed first. Although this hierarchical approach requires more complexity than the sequential approach, it does not result in any orphaned nodes in the service model instance 40.

SERVICE MODEL INSTANCE SPECIFICATION

Table 5 is an example of a portion of a service model instance specification for the service model template of Table 4 and the discovered instance of Table 3. The variables of the table are consistent with the variables of Tables 2-4. Referring now to FIG. 3, the service model instance 40 may be used by the view generator 42 to provide any of the three views. Preferably, the service model instance is represented as a graph of the nodes and edges that identify dependencies among the nodes. The service model instance is also used by the measurement agent configurator. Information from the

18

instance may be extracted by the configurator to merge test information relevant to particular elements with default algorithm descriptions in order to generate a measurement agent specification 46 for the ISP of interest.

TABLE 5

```

[10.174.173.23-fakeisp.net:DNS]
ipAddress=10.174.173.23
hostName=dns1.fakeisp.net
domainName=fakeisp.net
category=ExternalDnsServer
[10.174.173.23:SM-Host]
measurements=TCP-Connection-Rate(10.174.173.23),
VMstat(10.174.173.23), IFstat(10.174.173.23)
ipAddress=10.174.173.23
hostName=dns1.fakeisp.net
state=Alive
category=
[10.137.196.52:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.52),
VMstat(10.137.196.52), IFstat(10.137.196.52)
ipAddress=10.137.196.52
hostName=mailfes21.fakeisp.net
state=Alive
category=
[10.137.196.54:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.54),
VMstat(10.137.196.54), IFstat(10.137.196.54)
ipAddress=10.137.196.54
hostName=mailfes22.fakeisp.net
state=Alive
category=
[10.137.196.56:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.56),
VMstat(10.137.196.56), IFstat(10.137.196.56)
ipAddress=10.137.196.56
hostName=mailfes23.fakeisp.net
state=Alive
category=
[10.137.196.58:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.58),
VMstat(10.137.196.58), IFstat(10.137.196.58)
ipAddress=10.137.196.58
hostName=www.fakeisp.net
state=Alive
category=
[10.137.196.69:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.69),
VMstat(10.137.196.69), IFstat(10.137.196.69)
ipAddress=10.137.196.69
hostName=www.xyz.com.fakeisp.net
state=Alive
category=
[10.137.196.70:SM-Host]
measurements=TCP-Connection-Rate(10.137.196.70),
VMstat(10.137.196.70), IFstat(10.137.196.70)
ipAddress=10.137.196.70
hostName=www.abc.com.fakeisp.net
state=Alive
category=

```

CUSTOMIZING THE SERVICE MODEL TO AN ISP'S ORGANIZATIONAL STRUCTURE

As previously noted, the service model instance 40 of FIG. 4 is customized to an ISP's organizational structure, so that ISP operations personnel only view the status of services and service elements that are of relevance to the personnel. A straightforward approach to customizing the service model instance to an ISP's organizational structure is to edit the service model template and explicitly include nodes that capture the organizational dependencies. For example, the nodes may be grouped according to categories (e.g., InternalWeb services and WebHosting services). Each of these categories may be managed by different operations personnel. To accommodate this case, the service model template could be modified to define nodes that represent the

individual categories and dependencies that indicate the components of the different categories. Since the organizational structure varies from one ISP to another, the approach would require that each ISP edit the service model templates to match their organizational structure. Editing the service model template to define the categories and the components relationships can be a tedious task, especially for a large ISP.

An alternative approach is to allow an ISP to specify its organizational structure using the configuration interface 60 previously described with reference to FIG. 4. The service model template 34 is pre-specified to exploit the configuration specification and to generate a service model instance that is customized to each ISP. The main advantage of this approach is that the ISP operator only has to primarily edit the configuration specification, which is much less complex than editing the service model template 34.

The application of this less complex approach can be described with reference to Tables 1-5. Through the configuration interface 60, an ISP operator specifies ways in which the discovered services and service models are to be categorized in the discovered instance 36. In Table 1, the configuration specification indicates that the ISP uses the naming pattern *.com.fakeisp.net to identify web servers that are hosted for the businesses. Web servers that do not match this pattern are internal web servers that are used by the ISP's residential customers. Each of the discovery modules 50, 52 and 54 then uses the configuration specification to determine a categorization of the services and service models that are discovered. A discovery module is also included in the discovery template 48 to discover and report on all the categories that have been discovered in the ISP's system. This information is used by the service model creation engine 38 to construct a service model instance 40 that represents the ISP's organizational structure. To enable this, the service model template of Table 4 has a section that generates a node in the service model instance for each category in the discovered instance. The components list in this section maps all the services and service elements that are associated with a category to the corresponding node in the service model instance. Thus, by merely specifying the categorization of services and service elements using the configuration specification, an ISP operator can derive a service model instance that is customized for the ISP.

EXAMPLE: READ MAIL SERVICE MODEL

As an example implementation of the system of FIG. 3, the email service of "Read Mail" will be considered. The Read Mail service refers to a subscriber accessing his/her mailbox from the mail system of the ISP. FIG. 5 illustrates a service topology for this service. Using a client application that supports the Post Office Protocol-Version 3 (POP3), a subscriber at a desktop computer 62 attempts to access mail. Internal to the ISP system, the request from the subscriber's computer 62 may be received and processed by one of many servers 64, 66 and 68 that constitute a mail service group 70. The servers within the group are front-end servers (FESs).

Before the subscriber can access the appropriate mailbox, the mail server 64-68 that handles the request contacts an authentication server 72 to verify the identity of the subscriber. Typically, password identification is used in the Read Mail service. A subscriber database 74 is accessed in this process. Following the authentication process, the mail FES 64-68 accesses a mailbox 78 of the subscriber from a back-end content server 76 using the NFS service. The retrieved mail messages are transmitted to the computer 62 of the subscriber.

There are several active and passive measurements that can be made to assess the health of the different elements involved in supporting the Read Mail service. A measurement system (MS) may be installed in the server farm of the ISP to perform measurements using agents executing on the MS and on the different ISP hosts. The different measurements that characterize the Read Mail service include an active service quality measurement of availability and response time made from the MS in the service farm, an active measurement of network throughput from the MS in the service farm to the POP sites, passive measurements of CPU and memory utilization passive measurements of TCP connection traffic and packet traffic to the mail servers obtained from agents executing on the mail servers, and passive measurements of NFS statistics (e.g., number of calls, timeouts, and duplicate transmissions) on the mail servers and the mail content servers. The active measurements attempt to assess the service quality as viewed from subscribers connecting to the POP sites, while the passive measurements may be used to assess resource utilization and traffic statistics that are critical for problem diagnosis.

FIG. 6 is an illustration of an example of a view that may be presented to an ISP operator using the view generator 42 of FIG. 3. While the oval-shaped nodes in the service model graph represent the different services and service elements, the arrows represent measurements of services and service elements. The root of the service model graph is the Read Mail service, represented by oval 80. The state of this node represents the overall health of the Read Mail service, as assessed by the MS located in the service farm of the ISP. That is, the overall health is assessed without considering the state of the network links from the server farm to the POP sites. In one embodiment, the overall health is represented by color coding the oval 80. For example, oval 80 may be shaded green to designate a positive health of the Read Mail service, and may be shaded red if the Read Mail service has degraded in its availability or performance.

Typically, there is no direct measure of the overall health of the Read Mail service. Instead, the state of the service must be inferred, based on the states of the different mail FESs 64-68 that together enable the Read Mail service. Direct active measures of availability and performance, and passive measurements of TCP statistics to the POP3 service port, together contribute to the determination of the status of the Read Mail service. The active and passive measurements are performed at each of the mail FESs, as indicated by the arrows corresponding to the second level service oval 82 in FIG. 6.

The next level of the service model graph reflects the dependencies of the Read Mail service on one element and two services. Oval 84 is the dependency of the Read Mail service on a POP3 server executing on the mail FES. Oval 86 represents the authentication service for verifying the identity of the subscriber from which a Read Mail request is received. Oval 88 represents the NFS service used by the particular mail FES. Considering the POP3 server 84 first, the health of the server is measured based on the ability to establish a TCP connection as part of the active Read Mail service quality measurement and the time required to establish the connection. In turn, the health of the POP3 server may be impacted by the link, represented by oval 90, interconnecting the mail FES to the measurement station (from which the active test is run) and the health of the mail FES host, represented by oval 92. As shown in FIG. 6, both the link 90 and the host 92 include four performance parameters that are measurable in determining the health of the nodes. While not shown in FIG. 6, the state of the various

nodes may be represented by color coding or by other display means for distinguishing the states of the nodes.

Regarding the dependency of the authentication service 86 on the mail service 82, since it is possible that the authentication service is healthy but a specific mail FES is failing to perform authentication, the authentication service is first represented from the point of view of each of the mail FESs 94. Direct measures of the authentication delay when accessing through a mail FES are used to determine the state of the mail authentication service 86 in relation to that mail FES 94. The service model for an authentication service node, whose state affects the state of the mail FES-specific authentication node, is not expanded in the service model graph of FIG. 6. Likewise, the NFS service 88 is not shown as being expanded in the service model graph. The service model dependency is handled in much the same way as the authentication service dependency.

As previously noted, the service model graph of FIG. 6 represents the Read Mail service in isolation. To represent the end-to-end service, a service model must take into account the state of the DNS service used by subscribers to resolve the ISP mail domain to each one of the mail FESs, and the state of the network links between the different POP sites and the ISP server farm. Clearly, since the different POP sites use different routes to connect to the service farm, a subscriber's perception of the end-to-end service may vary, depending upon the POP site that the subscriber is using. FIG. 7 is a service model graph for the Read Mail service as perceived by a subscriber connected to the ISP system via POP_m. The Read Mail POP_m service is represented by oval 96 and has the Read Mail service 80 of FIG. 6 as one of its dependencies. However, the service 80 is not expanded. While not shown in FIG. 7, the graphing preferably includes color coding or other designation for nodes, such as the service 80, which are not fully expanded.

The other dependencies on the Read Mail POP_m service 96 include the link 98 and the DNS service 100. The health of the link is determined by measurements of the performance parameters throughput, packet loss, delay and connectivity. The health of the DNS service 100 is determined by measurements of availability and performance. The DNS service 100 is shown as having the dependency POP_m DNS server 102. In turn, the server 102 has two dependencies, namely POP_m DNS host 104 and the link 106.

There are several ISP management functions that can exploit the capabilities of the service models. For example, a service model for operational monitoring may be configured to indicate the status of the different elements providing a service. When a failure occurs, the service model indicates which element or elements have been affected by the failure. Moreover, by traversing the service model graph top-down, an operator can determine the root-cause of the failure. For example, in FIG. 6, when a problem is noticed with the overall Read Mail service 80, an operator can traverse the service model graph to determine whether the problem is caused by a specific mail FES or whether all of the mail FESs are experiencing a problem. Assuming a similar scenario, moving down the service model graph, the operator can further determine whether the problem is related to authentication failure, NFS failure, or a failure of the POP application server 84.

Since services and service elements can be organized based on domains of responsibility in a service model, ISP operations personnel need only monitor and diagnose the services that fall in their domain of responsibility. In the Read Mail service example of FIG. 5, an email operations

expert who is responsible for the mail service and the mail servers uses the service model depicted in FIG. 6, since the expert is mainly interested in the states of the email services and servers. The authentication and NFS services are included in the service model representation, since these services can adversely impact the Read Mail service. In contrast, the links between the service farm and the POP sites are not included in the model, since they do not affect the Read Mail service from the perspective of the email expert.

MEASUREMENT TOPOLOGY AND STATE REPRESENTATION

As can be seen in FIGS. 5 and 6, the service model maps different measurements for some of the nodes in the service model graphs. The node to which a measurement maps depends on the semantics of the measurement (i.e., which logic node or nodes are targeted by the measurement) and the location or locations from which the measurement is made. In the simplest case, each measurement directly maps to one of the nodes in the service model. In some cases, measurements may span multiple service elements and there may not be a direct mapping of a measurement to a node in the service model. In such cases, composite measurements, which are combinations of measurements being made in the ISP system, may have to be composed and mapped to the nodes in the service model. For example, suppose Link (x,y) is a network link interconnecting hosts x and y in an ISP system, and suppose Link (x,y) is comprised of Link (x,z) and Link (z,y). If measurements are made from x, Link (x,y) and Link (x,z) can be directly measured. The status of Link (z,y) has to be derived from the status of Link (x,y) and Link (x,z).

Each of the nodes in the service model has an instantaneous state associated with it. As previously noted, the state of a node reflects the health of the service or service element that it represents. A number of policies can be used to compute the state of service model nodes. One policy computes the state of a node based on all of its measurements. Another policy assigns weights to the different measurements based on an estimate of their reliability, as gauged by a domain expert. Yet another policy determines the state of a node based on its measurements in combination with the states of its dependencies in the service model graph. The states of the measurements associated with a node may be determined by applying baselining and thresholding techniques to the measurement results.

DISCOVERY METHODOLOGIES

Returning to FIGS. 2 and 4, the preferred embodiment of the management system 56 includes a discovery engine 58 that automatically discovers services, service elements and dependencies among services and service elements. FIGS. 7 and 8 illustrate applicable discovery methodologies. Similar to existing management system implementations, a first phase of discovery is performed from a management station 108, which may be external to the service farm of the ISP system. Predominantly, this phase involves active tests that generate test traffic and query all of the ISP hosts to detect the existence of different types of servers 12, 14, 22, 26 and 28. That is, the phase detects execution dependencies, as defined above. Component and organizational dependencies are also detected during this phase. Moreover, some of the inter-service dependencies are discovered, but the second phase is focused on the inter-service dependencies, since it may not be possible to discover all the inter-service depen-

dencies that exist using tests executed from outside the ISP host. The second phase uses an internal view of the ISP system. Preferably, the two phases are executed sequentially, with the second phase utilizing the discovered information output by the first phase to direct its operations. Different mechanisms can be employed in the internal discovery phase. Both phases of discovery must be executed periodically, so as to discover new services and service elements that may have been introduced into the ISP environment.

In FIG. 8, a single discovery agent 110 is used in the first phase external discovery process. The solid line 112 represents the discovery agent contacting the DNS server 14 to get a list of hosts in the ISP system. The dashed lines from the discovery agent 110 indicate active tests being executed by the discovery agent to detect various types of dependencies that exist.

FIG. 9 is an illustration of the second phase internal discovery process. In this phase, a number of internal discovery agents 114, 116, 118 and 120 are utilized. The solid lines having arrows at one end indicate the discovery of inter-service dependencies. The dashed lines indicate the flow of discovered instance information back to the management station 108.

Both phases of the discovery process may employ one or more different basic categories of discovery techniques. Under the service-independent techniques, hosts and networking equipment (routers 20, hubs 18, and POP sites 16) which are not specific to any service are discovered. As a second category, service-generic techniques (which may be the same techniques, but with appropriate parameterization for each service) may be used to discover instances of different services. In order to do so, typically, such discovery techniques exploit common characteristics of different services to discover instances of the services. An example of this second category is a discovery technique for News, Web and email services. Since all of these services rely upon the same transport protocol (i.e., TCP), a discovery technique can discover the existence of application servers by connecting to the TCP ports corresponding to the services.

Another category of techniques may be referred to as the service-specific-but-application-independent techniques. Techniques in this category are specific to the service. They are intended to monitor, but may be used for discovery that is independent of the specific application server that is being used to implement the service. For example, the discovery of the relationship between the services offered by POP3 email servers and the service provided by SMTP-based email servers is possible using application-independent techniques, since the relationship is accessible from the domain name service in the form of mail exchange (MX) records.

A fourth category may be referred to as application-specific techniques. Many inter-service dependencies may need to be discovered in a manner that is very specific to the application servers providing the services. For example, to discover the dependency of the web service on NFS, the discovery technique may have to query the configuration file in the web application server that is providing the web service. Since the format and contents of a web application server's configuration file are specific to the application, the discovery technique is application-specific.

FIRST PHASE EXTERNAL DISCOVERY

An often under utilized component of an ISP system is the domain name service (DNS). In order to allow subscribers

to access hosts using their host names, rather than their more difficult to remember IP addresses, DNS stores the host name-to-IP address mapping for all of the hosts in the ISP system. Moreover, the exchange of email messages across hosts occurs using the mail exchange records maintained by the DNS. Name server (NS) records in the DNS database serve to identify authoritative name service for the ISP domain—these are name servers that are externally accessible from the global Internet and are the authorities that are contacted when users in the global Internet attempt to access any hosts in the ISP system. Moreover, service groups, such as an email service group, are enabled via round-robin scheduling mechanisms implemented in the DNS servers. In summary, the domain name system holds a wealth of information that is critical for auto-discovery of ISP services. However, additional mechanisms are necessary to complement DNS-based discovery mechanisms.

One of the first steps in discovery is to determine all of the hosts that exist in the ISP system. Most existing network management systems have taken one of two approaches. A first approach is to scan an address range that is either specified by an operator or is determined based on the local subnet of the measurement host. The address range is scanned using ping to solicit responses from all IP-enabled hosts. The second approach is to use the default router configured for a host to boot-strap the discovery. Communication using SNMP with the router and using its routing tables may be used to discover hosts in the local subnet. Also, routing table information may be used to discover other routers that can provide additional host information.

An alternative approach that is complementary to either of the two traditional approaches is to obtain a list of all of the hosts in the ISP system using information available with the domain name service. From the host name of the management station in which the discovery process executes, the ISP domain name can be deduced. Using the default name service configured for the management system, the discovery process queries DNS to obtain the NS records that list the authoritative name servers for the ISP domain. One or more of these name servers are contacted by the discovery process to obtain a list of named hosts in the ISP system. Zone transfer capabilities supported by all DNS servers are used for this purpose, as is known in the art. While ISPs usually manage a single domain, some ISPs may have multiple domains under their control. To discover all of the hosts that exist in the different domains, the discovery process must be informed of the existence of all the different domain names. This information cannot be automatically deduced by the discovery process.

The steps that are executed in the first phase of discovery are identified in FIG. 10. It is not critical that the steps be followed in the order shown in the figure. Following the step 122 of discovering the hosts, the application servers are discovered in step 124. The existence of application servers of different types (e.g., Web, Email, News, FTP, DNS, NFS, and Radius) is verified by active tests that emulate typical client requests directed at the different TCP and UDP ports corresponding to the different service types. A response to an emulated request has to be interpreted in a service-specific manner. By observing the header in a response from the web service, the discovery process determines the type of web server that is executing on the host machine. Likewise, by processing the response returned by the email and News servers, the discovery process determines the type of servers. This information can be used to customize the second phase of discovery. For instance, discovery agents installed on the ISP host machines in the second phase of discovery

may process a web application server's configuration files to discover NFS dependencies that the server may have. Since web server configuration files are typically specific to the type of server, the server type information provided by the first phase of discovery can be used to determine the processing capabilities of the discovery agent or agents that must be deployed in the second phase. The server type information may also be used to determine specific measurements which must be targeted for the server to monitor its status.

In step 126, the existence of Web, Email, News, DNS and other service groups has to be determined using the DNS. By querying the DNS, the discovery process determines a list of domain names that have multiple IP addresses associated with them. For each name in the list, the discovery process then determines whether each of its IP addresses hosts a common application server. If so, the name likely represents a DNS round-robin service group that supports a common service. For example, suppose that all of the IP addresses corresponding to the name www.fakeisp.net host web servers. In this case, www.fakeisp.net represents a web service group. Note that in this process, a host that has two network interfaces, and therefore is assigned to different IP addresses, may be listed as a service group. Using the virtual host/server discovery heuristics discussed below, all such hosts can be removed from the service group list.

In step 128, the MX records for the ISP domain are accessed from the DNS system. The MX records indicate the mail servers that must be contacted to deliver mail to the ISP domain. The list of SMTP-based mail servers thus determined represent the servers that handle delivery of incoming mail to the subscribers of the ISP. Discovery of these servers is essential to automatically generating a service model for the email service that represents the delivery of mail from the Internet to the subscribers. Moreover, the mail gateways may be managed by a different entity than the one that manages mail servers that are internal to the ISP.

One of the critical measures of the performance of an email service of an ISP is the round-trip delay between the transmission of an email message from a source host and its reception at the intended destination. This measurement can be used to assess email delivery times in the ISP domain, from the ISP domain to locations on the Internet, and from locations on the Internet to the ISP domain. Since the email service uses different protocols and, hence, different application servers to send mail and to receive mail, in order to initiate round-trip delay measurements of email, it is essential to determine relationships between the different types of email servers on the ISP domain (e.g., which SMTP server can be used to send mail to a POP3/IMAP-based server). This discovery is executed at step 130. Since mail forwarding is predominantly based on the MX records maintained in the DNS database, by querying the DNS system for MX records corresponding to each of the POP3/IMAP servers, the discovery process determines the mail service relationships in the ISP domain.

Various approaches can be adopted to discover the terminal servers that exist in an ISP POP site in step 132. The most straightforward approach uses SNMP queries to obtain the MIB-II system description from all the hosts in the ISP network. Based on the replies, the discovery process can identify the hosts that are terminal servers. An alternative approach is based on the observation that because they need to operate and manage thousands of terminal servers, most ISPs have specific naming conventions that they use when naming their terminal servers. In fact, the naming convention more often indicates the association between terminal

servers and POP sites, so that when a problem is reported by a subscriber using a POP site, the ISP operations staff can quickly decide which of the terminal servers needs to be checked in order to diagnose the problem. With this approach, an ISP provides a regular expression representing the naming convention used as input to the discovery process. By matching the list of hosts that are discovered with the naming convention, the discovery process not only determines the terminal servers that exist, but also determines the POP site to which a terminal server is assigned. Another key advantage of this approach is that it performs discovery without generating any additional network traffic.

The approach of exploiting name conventions may also be used in step 134 to categorize the other services of the ISP. As an example, for each web site that it hosts for its business customers, a particular ISP may assign internal names of the form *.com.fakeisp.net, so that a hosted web site named www.customer-domain.com will have a corresponding entry for www.customer-domain.com.fakeisp.net in the internal DNS database of the ISP. As in the case of terminal servers, by permitting an ISP to specify its naming conventions, the discovery process composes a categorization of services that is customized to the target ISP system. This categorization can be based on geographical locations of services, based on business relationships, and/or based on the delegation of responsibilities among operators. The categorization information can be used to automatically define the customized service model for each ISP, with special nodes in the service model representing a collection of nodes pertaining to the same category.

SECOND PHASE DISCOVERY

By treating the ISP system as a "black box," the first phase of external discovery detects most of the execution, component and organizational dependencies of the ISP. Additionally, some of the inter-service dependencies are discovered. The internal discovery process is focused solely on detecting inter-service dependencies, particularly those that are not discovered by taking an external viewpoint. For example, the relationship between a mail server and an NFS server is not discoverable from an external viewpoint.

There are two basic approaches for conducting the internal discovery. One approach uses network probes, while the other approach uses special-purpose discovery agents. Regarding the first approach, software probes installed at strategic locations on the network can snoop on packet transmissions. Since most TCP/IP communication is based on source/destination port numbers, by processing the headers of packets that are captured by the probe, a software probe can deduce many of the relationships that exist among services. For example, a UDP packet transmitted from a mail server to the NFS port of an NFS server indicates that the mail server depends on the NFS server for its content storage.

An advantage of the approach that utilizes network probes is that the approach enables discovery of inter-service dependencies independent of the specific types of application servers residing on the host. Moreover, since it relies on just the ability to capture packets on the wire, this approach handles UDP and TCP-based services equally well.

The key difference between the approach of using network probes and the approach of using special-purpose discovery agents is that unlike the network probes, the discovery agents do not snoop on packet transmissions. Instead, the discovery agents use a number of operating systems and application-specific mechanisms to discover

inter-service dependencies. These mechanisms include (1) processing service configuration information and (2) application-dependent monitoring tools. Referring first to the processing service configuration information, application servers determine their dependencies on other services from one or more configuration files. By processing the content of the configuration files, discovery agents can discover inter-service dependencies. An example of this is the processing of the web server's configuration file to discover whether it has a dependency on an NFS service. While processing the web server's configuration file, the discovery agent can also determine if the same application is being used to host multiple "virtual" servers (which is commonly used by ISPs to host web sites on behalf of their business customers). Typically, web server configuration files are specific to the type of server executed on the web server in use. The server type determination performed during external discovery (i.e., the first phase of discovery) is used for deciding the location and format of the configuration files.

While many Unix operating systems use configuration files that are defined in an application-specific manner, Windows NT-based systems store all application configuration information in the registry. In the Windows NT systems, while the registry can be processed in an application-independent manner, the specific configuration attributes have to be interpreted in an application-specific manner.

Thus far, only forward-looking discovery agents have been identified. These are agents that discover dependencies of a service on other services by querying configuration files of the application providing the service. Sometimes it is easier to implement backward-looking discovery agents to discover the dependencies on a service (i.e., discover which other services are using the service). For example, the configuration file of the mail authentication server may indicate which of the mail servers are depending on the authentication server. One of the ways of implementing backward-looking discovery agents is by processing application server configuration files.

Turning now to the second mechanism of using application-independent monitoring tools, this approach is particularly attractive for services that communicate using TCP. The netstat utility can be used to determine the TCP connections that exist on an ISP host. A discovery agent that executes this tool can periodically discover information about the source and destination ports and the host locations for TCP connections, which can then be used to deduce inter-service dependencies.

This second approach of using application-independent monitoring tools exploits the fact that most TCP implementations enforce a three-minute delay for connections in the TIME_WAIT state of TCP, so that a connection persists for about three minutes even after it is no longer in use. Consequently, whenever the discovery agent is executed, it is likely to detect all the TCP connections that may have been established in the three minutes prior its execution. This same approach does not work for UDP-based services, since UDP is connectionless, and there is no state that is maintained at either the source or the destination.

The approach of monitoring TCP connections can be used to discover dependencies such as those that exist between mail servers and mail authentication servers, between servers and back-end databases, between Radius/TACACS authentication servers and terminal servers, and between similar relationships. Again, discovery agents can be forward-looking or backward-looking.

Advantages of using discovery agents, as compared to network probes, include the reduction of overhead on the ISP hosts, the relaxation of security concerns, and the fact that all of the discovery agents do not need to be deployed at the same time. Instead, the deployment of discovery agents can occur at the discretion of the ISP. As and when new discovery agents are installed on the ISP hosts, additional information is discovered about the ISP system.

FIG. 11 is a process flow of steps relevant to the second phase of the discovery process. In steps 136 and 137, the information that is obtained in the first phase of discovery is used to generate an incomplete service model instance. As previously noted, the first phase of discovery provides the necessary information for identifying component dependencies, organizational dependencies, execution dependencies and some of the inter-service dependencies. A first instance generator matches the service model template with the auto-discovered information from the first phase to generate the incomplete service model instance. However, other inter-service dependencies are not discoverable using the techniques of the first phase (e.g., the relationship between a mail server and an NFS server).

In steps 138 and 139, the holes in the incomplete service model instance are identified and information obtained in the first phase is used to determine appropriate discovery actions. The incomplete service model instance is used to determine the types of relationships that must be examined. In the mail service example, if a host is discovered in the first phase to be running a POP3 server, the second phase may be used to discover the name file service and authentication service used by the POP3 server on the particular host.

In step 140, the network probes and/or the special-purpose discovery agents are deployed in a manner determined during execution of the step 139. For example, application-specific knowledge may be used to parse configuration files or log files, or may be used to search a configuration registry for a particular server instance executing on a particular host. The network probes and/or discovery agents generate service dependency outputs in step 141. The outputs are used in a second instance generation to complete the service model instance in step 143.

EXTENSIBLE DISCOVERY ARCHITECTURE

Since new network services and service elements are being deployed at a rapid pace, it is important that the discovery methodologies be implemented in an extensible manner, allowing new discovery capabilities to be incrementally added to the management system. FIG. 12 depicts the extensible architecture for discovery components previously described with reference to FIG. 4. The discovery modules 50, 52 and 54 represent the logic used for discovery of different services and service elements. The discovery template 48 is the key to the extensibility of the auto-discovery architecture in the sense that it drives how discovery is performed. The template defines the different services and service elements that need to be discovered, and the specific discovery modules that can be used to discover these elements. The template also establishes the format of the outputs from the modules.

The discovery engine 58 drives the auto-discovery process. The discovery engine interprets the discovery template 48 and for each of the service or service element types specified in the template, the engine invokes the corresponding discovery module 50, 52 and 54 specified in the template. All of the discovery modules report the results of their execution back to the discovery engine. The discovery

template contains instructions for the discovery engine to process the output of the discovery modules and to record them as a discovered instance 36.

Some discovery modules may rely on the discovery results of other discovery modules. For example, a DNS round-robin service group discovery module for web services relies on identifying which hosts support web services, which is an output of the web service discovery module 50. To accommodate these relationships, as part of the interface that the discovery engine 58 exposes to the discovery modules, the engine provides ways for accessing and searching the instances discovered by other discovery modules.

In contrast to the discovery modules 50, 52 and 54, the discovery engine 58 is designed to be independent of the services that need to be discovered. Consequently, to discover new services or service elements, a user merely has to provide a discovery template specification or one or more discovery modules for each new element. By providing an alternate discovery module for a service that is already supported, a user can also enhance capabilities of an existing discovery system.

In practice, there are two significantly different approaches to designing the discovery engine 58 and the discovery modules 50-54. A first approach is to enable the discovery engine to control the discovery process. The discovery engine accesses the discovery template and determines the order in which sections of the template are processed. On the other hand, in the second approach, the discovery modules drive the discovery process. In effect, this is an event-driven approach, since the results obtained from one module will trigger subsequent activities by other modules.

Regarding the first approach in which the discovery process is driven by the discovery engine 58, FIG. 13 illustrates the logical building blocks of the discovery engine. The discovery engine executes periodically and each time it starts, the engine processes the discovery template.

By considering the values of the dependencies variable for each of the sections in the discovery template, the discovery engine determines the order in which the sections must be processed. Thus, the discovery engine includes a template parser 142. Sections of the template which have no dependencies are processed first. A module loader 144 directs the relevant information to the appropriate discovery module 146 for processing a particular section in which no dependencies are identified. After all such sections are processed, the discovery engine iterates through the list of template sections, choosing sections which have not been processed and which have their dependencies determined by earlier processing. This process is repeated periodically to discover new instances as and when they are added to the system being discovered. In one application, the discovery engine uses the exec system call to invoke the discovery modules at separate processes. By doing so, the discovery engine is able to handle discovery modules written in a variety programming environments.

A query processor 148 of the discovery engine 58 performs two functions. First, when a module 146 is activated, the processor 148 queries the discovery engine to obtain configuration information that guides the discovery modules. In FIG. 4, the configuration information is generated from the configuration interface 60 that is manipulable by a user. Table 1 was previously included to depict a typical configuration file. Each line in the file represents an instruction for one of the discovery modules. The first column of the line identifies the discovery module to which the instruc-

tion pertains. There are three types of instructions that are specified in the configuration file. All of these instructions specify regular expression patterns that must be applied against the IP address or host name of the service or service element. The instructions in the configuration file are (1) criteria that instruct the discovery modules to include or exclude specific services or service elements, (2) criteria that instruct the discovery modules to associate specific services or service elements with certain categories, and (3) criteria for discovering terminal servers and for extracting POP site-to-terminal server mapping from the terminal server names.

The second function of the query processor 148 is to provide the discovery modules 146 with access to previously discovered instances. Based on configuration and discovered instance information obtained from the query processor, the discovery modules perform tests on the ISP system and report their discovery output to the discovery engine. A discovery instance generator module 150 of the discovery engine processes the results of the discovery modules and outputs the discovery instance in an appropriate format. An example of such a format was previously set forth in Table 3. The formats of the discovery template and the discovery instance are thereby hidden from the discovery modules.

As previously noted, the second approach to designing the discovery engine 58 and the discovery modules 50-54 of FIG. 12 is to establish an arrangement in which the discovery process is driven by the modules. In this alternative embodiment, the discovery engine processes the template once, invoking the discovery modules simultaneously. From this point, the discovery modules determine when different elements in the ISP system are discovered. The discovery modules execute periodically, looking for new instances. Some discovery modules are independent in the sense that they are not reliant on other modules for discovery. These independent modules begin executing immediately.

As and when a discovery module 50-54 discovers a new instance, the discovery module forwards its results to the discovery engine 58. Based on the dependencies on a discovery module, as specified in the discovery template 48, the engine 58 forwards the results to other discovery modules for which the results are relevant. The availability of new results (e.g., the discovery of a new host) may trigger discovery by other modules (e.g., the web server module checks to determine if a web server is executing on the new host), and this process continues. A key advantage to this approach, as compared to the engine-driven discovery approach, is that multiple discovery modules may be executing in parallel, discovering the ISP's services. In this approach, the discovery engine 58 mainly functions as a facilitator of communication among the discovery modules. A variant of this approach may not even involve the discovery engine, with the discovery modules registering interest in other discovery modules and information concerning newly discovered instances being directly communicated among the discovery modules, without involving a discovery engine.

INTEGRATING DISCOVERY WITH SERVICE MODELS

In the scenario in which the management system uses service models for management of Internet services, there are two ways in which discovery can be integrated with service models. In a looser integration, the output of discovery (the discovered instance) is integrated with a service model template that outlines the structure of a service, and

the integration automatically generates a service model instance that is customized for the ISP system being managed. However, the preferred integration is one that provides a tighter integration, and involves driving auto-discovery and service model instantiation from a common template. In this preferred approach, for each node in the service model, corresponding discovery template specifications are provided. The discovery and service model-specific components of the template can either be processed in a single application or can be processed separately. This approach towards tighter integration of discovery and service model templates is attractive for several reasons. Firstly, the service model template can serve to constrain the discovery process, since only services and service elements that are specified in the service model template need to be discovered. Secondly, depending upon its design, the service model template could end up using some of the outputs of the discovery process. Using a common template permits tighter syntax checking across the discovery and service model components of a template. Thirdly, the two-phase approach to discovery described above fits in well with the service model concept. The inter-service dependencies that need to be discovered in the second phase (internal discovery) can be determined based on the service model template. Finally, the discovery process itself can be determined based on the service model template specification. The discovery process may attempt to traverse the service model template tree from a root node down. At each level, it attempts to discover all services or service elements of the types specified by the node, providing all of the children of a node have been discovered. If this is not the case, the discovery process proceeds to first discover all instances of the children nodes. Continuing the tree traversal recursively, the discovery process discovers all instances that are necessary to build the service model for the ISP system being managed.

What is claimed is:

1. A method of modeling services available via networks comprising steps of:

selecting a core service of interest;

forming a discovery template that is specific to said core service such that said discovery template includes instructions for implementing automated techniques for discovering instances of services and service elements which are anticipated as being cooperative in execution of said core service within a network environment, including formatting said discovery template for access during implementation of said automated techniques; and

displaying discovered instances of said services and service elements in isolation from said discovery template.

2. The method of claim 1 wherein said step of forming said discovery template includes identifying types of said service elements without specifying actual service elements, said step further including identifying types of services which are anticipated as being cooperative in execution of said core service.

3. The method of claim 2 wherein said step of identifying said types of service elements includes specifying types of servers and wherein identifying said types of services includes specifying services enabled by said types of servers.

4. The method of claim 2 wherein said step of forming said discovery template includes associating each identified type of service and type of service element with a discovery module for implementing said automated techniques.

5. The method of claim 4 wherein said step of forming said discovery template further includes specifying depen-

dencies among said discovery modules with respect to sequencing deployment of said discovery modules in implementing said automated techniques.

6. The method of claim 4 wherein said step of forming said discovery template includes providing instructions regarding organizing outputs of each said discovery module for display during said step of displaying discovered instances, said step of displaying being executable in the absence of said discovery template.

7. The method of claim 1 wherein said formatting of said discovery template includes storing said discovery template in a format compatible with a discovery engine for driving said automated techniques.

8. The method of claim 1 wherein said formatting of said discovery template includes arranging said instructions in an order which determines a sequence for implementing said automated techniques, said step of displaying including employing a service model template to determine a display format.

9. The method of claim 1 further comprising a step of enabling cooperation between memory for storing said discovery template and software for implementing said automated techniques such that said software determines a sequence for implementing said automated techniques, said step of displaying including employing a service model template to determine a display format.

10. A method of modeling a selected service available via a network comprising steps of:

enabling discovery capabilities for executing diverse discovery routines specific to acquiring decentralized information that identifies services and service elements of said network;

forming a discovery template that is specific to said selected service, including providing data relevant to triggering selected said discovery routines for acquiring said decentralized information, said selected discovery routines being related to acquiring instance information identifying services and service elements accessed in providing said selected service;

storing said discovery template; and

activating an automated process for collecting said decentralized information so as to drive modeling of said selected service, including accessing said discovery template to initiate said selected discovery routines.

11. The method of claim 10 wherein said step of forming said discovery template includes identifying types of service elements, including identifying at least one server type such that said step of activating said automated process includes transmitting queries that are specific to discovering instances of servers of said server type.

12. The method of claim 11 wherein said step of enabling said discovery capabilities includes providing software discovery modules specific to an Internet Services Provider (ISP) system, said discovery template being formed without being specific to said ISP system, thereby enabling use of said discovery template with any one of a plurality of ISP systems.

13. The method of claim 10 wherein said step of forming said discovery template includes providing instructions for configuring outputs generated in response to said discovery routines, said modeling being performed without entering said outputs into said discovery template.

14. The method of claim 10 wherein said step of forming said discovery template includes forming template sections such that each template section (1) is specific to a type of service element, (2) specifies at least one said discovery routine for identifying a service element of said type, and (3)

specifies dependencies that said at least one discovery routine has on outputs of other said discovery routines.

15. A system for modeling a selected core service available via a network comprising:

- a plurality of discovery module means for accessing decentralized information available on said network and for generating outputs indicative of service elements of said network;
- a discovery engine means connected to each of said discovery module means, said discovery module means being responsive to said discovery engine means with respect to invoking individual discovery module means; and
- a stored discovery template that is specific to said core service to be modeled, said discovery template including data indicative of selected said service elements which are anticipated as being cooperative in execution of said core service, said discovery template further including data indicative of selected discovery module means for identifying said selected service elements, said discovery engine means being connected to access said discovery template such that said discovery engine means automatically invokes said selected discovery module means;

wherein said discovery template controls said discovery engine means and wherein said discovery engine means controls said discovery module means, said outputs that are indicative of said service elements being recorded for representation of said core service independently of said discovery template.

16. The system of claim 15 wherein at least some of said discovery module means are configured to generate outputs

indicative of selected services which are anticipated as being cooperative in execution of said core service, said stored discovery template having data indicative of said at least some discovery module means.

17. The system of claim 16 further comprising a service model instance engine means for generating a graphical representation of said core service in response to said outputs of said discovery module means, said graphical representation including nodes representing said service elements that are indicated within said outputs of said selected discovery module means, said service model engine means being independent of said discovery template.

18. The system of claim 16 wherein said discovery template further includes information indicative of operational dependencies among said discovery module means in discovering said selected services and selected service elements.

19. The system of claim 18 wherein said discovery template has an arrangement that controls a sequence of invoking said selected discovery module means, said arrangement being at least partially determined by said dependencies.

20. The system of claim 18 wherein said discovery engine means determines a sequence of invoking said selected discovery module means, said sequence being at least partially determined by said dependencies.

21. The system of claim 18 wherein said discovery module means cooperate to determine a sequence of generating said outputs, said sequence being at least partially determined by said dependencies.

* * * * *